Computing Skylines on Distributed Data

Haoyu Zhang
Indiana University Bloomington
Bloomington,IN,USA
hz30@umail.iu.edu

Qin Zhang
Indiana University Bloomington
Bloomington,IN,USA
qzhangcs@indiana.edu

arXiv:1611.00423v1 [cs.DB] 1 Nov 2016

ABSTRACT

In this paper we study skyline queries in the distributed computational model, where we have s remote sites and a central coordinator (the query node); each site holds a piece of data, and the coordinator wants to compute the skyline of the union of the s datasets. The computation is in terms of rounds, and the goal is to minimize both the total communication cost and the round cost.

Viewing data objects as points in the Euclidean space, we consider both the horizontal data partition case where each site holds a subset of points, and the vertical data partition case where each site holds one coordinate of all the points. We give a set of algorithms that have provable theoretical guarantees, and complement them with information theoretical lower bounds. We also demonstrate the superiority of our algorithms over existing heuristics by an extensive set of experiments on both synthetic and real world datasets.

1. INTRODUCTION

Skyline computation, also known as the *maximal vector problem*, is a useful database query for multi-criteria decision making. If we view data objects as points in the Euclidean space, then the skyline is defined to be the subset of points that cannot be dominated by others, where we say a point y dominates a point x if y dominates x in all dimensions. This problem was first studied in computational geometry in the mid-1970's [10], and was later introduced into databases as a query operator [4]. There is a vast literature on skyline computation and its variants, and we refer readers to [6] and [8] for excellent surveys.

Most work on skyline computation in the literature has been conducted in the RAM (signal machine) model. In recently years, due to the large size of the datasets and the popularity of *map-reduce* type of computation, a number of parallel skyline algorithms have been proposed [19, 15, 18, 17, 9, 1, 13, 14]. A common feature of these parallel algorithms is that they use the *divide-and-conquer* approach, that is, they use central mechanisms to partition the whole point set into a number of subsets, and then assign each subset to a machine for local processing; finally the local results are merged to form the global skyline. The art of the algorithm design in this line of work lies in how to choose the partition mechanism.

In this paper we study the skyline computation on distributed data, which is different from parallel computation in that the data is inherently distributed in different locations, and we *cannot* afford to repartition the whole dataset since data repartition is communication prohibited over networks, and may also cause local storage and data privacy issues which the query node cannot control.

Consider a global hotel search engine, where each hotel is represented as a point in the 2-dimensional Euclidean space with the x-coordinate standing for the price and the y-coordinate standing for the rate of the location. A user naturally wants to find a hotel with the best location *and* the best price, although in reality hotels in good locations usually have higher prices. Thus a good search engine should recommend the user with a list of candidates such that no other hotel has both cheaper price and better location. This list is exactly the skyline of the point set. Given a query, the search engine (e.g., *kayak.com*, *hipmunk.com*) needs to contact servers/providers in different locations worldwide. The total bits of communication between the query node and servers and the communication rounds typically dominate the engine's response time, since sending messages through network is much slower than local computation, and the initialization of a new communication round takes quite some system overhead.

The Coordinator Model and Previous Work. We study the skyline problem in the *coordinator model* which captures the type of distributed computation mentioned above. In this model we have s remote sites each holding a piece of data, and a central coordinator which acts as the query node. We assume there is a two-way communication channel between each site and the coordinator. The computation is in terms of rounds: at the beginning of each round the coordinator sends a message to some of the sites, and then each of the contacted site sends a response back to the coordinator. The goal is to minimize the total bits of communication and the number of the rounds of the computation. See Figure 1 for a visualization of the model.

We differentiate two scenarios of data storage at sites: *horizontal partition* and *vertical partition*. In the former each site contains an (arbitrary) subset of points. And in the later each site contains *one* attribute of *all* points together with their IDs which are used to restore the whole

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

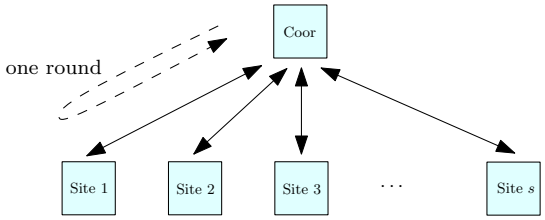

Figure 1: The Coordinator Model

point vectors. Thus if points have d dimensions, we have d sites. The vertical partition corresponds to the settings where the information of each attribute has to be retrieved from a different server/provider. For example, we may want to find the lowest prices of hotels from *kayak.com*, while the customers' best ratings from *TripAdvisor*.

The coordinator model is equivalent to the models used in several previous works for distributed skyline computation [15, 20] (horizontal partition) and [2, 12] (vertical partition). However, all the existing skyline algorithms in this model are heuristic in nature. We will briefly describe these algorithms in Section 4.1 (horizontal partition) and Section 3.2 (vertical partition) respectively. In [16] the authors studied a general version of vertical partition where multiple columns can be stored in one site. In this paper we will give algorithms with provable theoretical guarantees (for vertical partition we give a better heuristic, given a strong impossibility result that we will show). We will also compare our algorithms with previous ones experimentally in Section 4.

The Skyline Problem. We now give the formal definition of the skyline problem. In this paper we consider the problem in the 2-dimensional Euclidean space. Given two distinct points $u = (x_u, y_u)$ and $v = (x_v, y_v)$, we say u *dominates* v, denoted by $u \succeq v$, if $x_u \geq x_v$ *and* $y_u \geq y_v$. For a set of distinct points S, the skyline of S is defined to be

$$sk(S) = \{u \in S \mid v \in S, \text{ if } v \succeq u \text{ then } v = u\}.$$

We next define the data partition of the skyline problem in the coordinator model. In the horizontal partition case, each site i holds a set $S_i \subseteq S$. In the vertical partition case, we have two sites; the first site holds $\{(x_u, \mathsf{ID}_u) \mid u \in S\}$ where ID_u denotes the ID (key) of point u, and the second site holds $\{(\mathsf{ID}_u, y_u) \mid u \in S\}$. In both cases, the coordinator needs to output $sk(S)$ at the end of the computation.

Our Contribution. We have made the following contributions in this paper. Let n be the total number of points, k be the number of skyline points (i.e., the output size), and s be the number of sites.

1. For the horizontal partition, we propose two algorithms. The first one achieves the theoretically *optimal* communication cost $\Theta(ks)$, but needs $\lceil k/2 \rceil$ communication rounds. The second one gives a tradeoff between the total communication cost and rounds: given a round budget r the algorithm uses roughly $O(skn^{2/r})$ communication (see Theorem 2 for the precise bound). We also prove that any one-round algorithm needs linear (i.e., $\Omega(n)$) communication in the worst case. These are presented in Section 2.

2. For the vertical partition, we first prove that in the worst case, any algorithm needs linear communication, regardless of the communication rounds. We then propose a heuristic that works well in practical settings. These are presented in Section 3.

3. We have implemented our algorithms and relevant heuristic algorithms in the literature, and run them on both synthetic and real-world datasets. Our experiments have demonstrated the superiority of our algorithms over the existing ones in various aspects. We also noticed that for the horizontal partition and the datasets we have tested, with only three communication rounds the tradeoff algorithm can achieve similar communication cost compared with the theoretically optimal one. The experimental results are presented in Section 4.

Preliminaries. Let $[n]$ denote $\{1, 2, \ldots, n\}$.

The ϕ-quantile of a set S is an element x such that at most $\phi |S|$ elements of S are smaller than x and at most $(1-\phi)|S|$ elements of S are greater than x. If an ϵ-approximation is allowed (denoted by (ϕ, ϵ)-quantile), then we can return any ϕ'-quantile of S such that $\phi - \epsilon \leq \phi' \leq \phi$.

2. HORIZONTAL PARTITION

In this section we give two algorithms for skyline computation in the horizontal partition case. We then complement them by proving that the first algorithm is *optimal* in terms of the total communication cost (though it may need a large number of rounds). On the other extreme, if we want the computation to be done in one round then in the worst case the sites need to send almost everything to the coordinator. Our second algorithm gives a tradeoff in between.

We also show that if data is sorted among the sites according to one of the coordinates, then there is a simple algorithm that has much smaller communication cost and uses only two rounds.

2.1 Algorithms

There is a simple algorithm for computing the skyline points in the coordinator model in *one* round: Each site computes the skyline of its local data points and sends it to the coordinator, and then the coordinator computes the global skyline on top of the s local skylines. Unfortunately this algorithm has communication cost $\Omega(n)$; in other words, in the worse case almost all points in all sites need to be sent to the coordinator. We enclose a proof for this statement in Section 2.2.2. We thus try to explore if more rounds can help to reduce the communication cost.

2.1.1 Algorithm with Optimal Communication Cost

Our algorithm with communication cost $O(ks)$ is described in Algorithm 1. We will show in Section 2.2.1 that this communication cost is in fact optimal[1] even if we allow an *infinite* number of communication rounds.

Let us explain Algorithm 1 in words. At the beginning, each site computes its local skyline points since only these points can possibly be the global skyline points. The rest of the algorithm works as follows. At each round, the coordinator tries to find the (at most two) points with the largest x-coordinate or y-coordinate in the remaining points held by all sites. This is done by asking all sites to report

[1]Up to a logarithmic factor which counts the number of bits used to represent a point.

Algorithm 1 Optimal Communication under Horizontal Partition

Input: S_i is initialized as the point set held by Site i
Output: the global skyline
1: Site i computes its local skyline points and discards the other points in S_i
2: **while** $\exists i \in [s]$ *s.t.* $S_i \neq \emptyset$ **do**
3: **for** each i *s.t.* $S_i \neq \emptyset$ **do**
4: Site i sends the point with the largest x-coordinate and the point with the largest y-coordinate to the coordinator ▷ the two points can be the same point
5: **end for**
6: The coordinator computes new global skyline points from the points received from all sites, and sends the new global skyline points to each site
7: **for** each i *s.t.* $S_i \neq \emptyset$ **do**
8: Site i prunes S_i by new global skyline points received from the coordinator
9: **end for**
10: **end while**

the local maximums of their remaining points (Line 3-5). Next, the coordinator computes new global skyline points from the received local maximums, and then sends the new global skyline points to all sites for another local pruning step (Line 6-9).

We now show the correctness of Algorithm 1 and analyze its costs. First, it is clear that the points with the largest x-coordinate or y-coordinate must be on the skyline. After the pruning, the points with the largest x-coordinate or y-coordinate in the remaining points also must be on the skyline since they cannot be dominated by the other remaining points as well as the previous skyline points. We thus can find one or two skyline points in each round (finding one skyline point may only happen in the last round). Since there are at most k skyline points, the algorithm will terminate after at most $\lceil k/2 \rceil$ rounds.

The running time at each site consists of two parts: the computation of the local skylines and the point prunings. The computation of local skyline at the i-th site needs $O(n_i \log n_i)$ time [10], where $n_i = |S_i|$ is the number of points at the i-th site. The time used for pruning is linear in n_i at the i-th site: when pruning points using the new skyline point with the global maximum y-coordinate (Line 8), we scan the points which are sorted increasingly according to their x-coordinates after the local skyline computation. If a point cannot be pruned in a particular round, then all the points after it with larger x-coordinates cannot be pruned in this round. Therefore every point will only be pruned once in the whole computation. Same arguments hold for the prunings using the new skyline point with the global maximum x-coordinate. At the coordinator, for each round we only need to compute the maximums over at most $2s$ points. Thus the total running time is bounded by $O(ks)$.

THEOREM 1. *There exists an algorithm for computing the skyline on n points in the 2-dimensional Euclidean space in the coordinator model with s sites and horizontal data partition that uses $O(ks)$ communication and $\lceil k/2 \rceil$ rounds, where k is the output size, that is, the number of points in the skyline. The total running time at the i-th site is $O(n_i \log n_i)$ where n_i is the number of points at the i-th site, and that at the coordinator is $O(ks)$.*

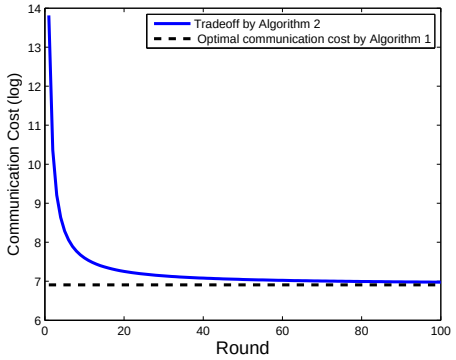

Figure 2: Theoretical Communication-Round Tradeoff

2.1.2 A Communication-Round Tradeoff

In the previous section we have shown an algorithm with the optimal communication cost but needs $\lceil k/2 \rceil$ rounds; on the other hand, there is a naive one-round algorithm but in the worst case it needs $\Omega(n)$ bits of communication. The natural questions is:

> *Can we obtain a communication-round tradeoff to bridge the two extremes?*

We try to address this question by proposing an algorithm that allows the users to choose the number of the rounds of the communication in the computation. In this section we show the following result.

THEOREM 2. *There exists an algorithm for computing the skyline on n points in the 2-dimensional Euclidean space in the coordinator model with s sites and horizontal data partition that uses r (≥ 3) rounds and $C = O(sk(n/s)^{1/\lceil r/2 \rceil})$ communication, where k is the output size, that is, the number of points in the skyline. The total running time at the i-th site is $O(C/s + n_i \log n_i)$ where n_i is the number of points at the i-th site, and the total running time at the coordinator is $O(C)$.*

Figure 2 visualizes the communication-round tradeoff.

Let $t = \lceil r/2 \rceil$. We describe our tradeoff algorithm in Algorithm 2. The algorithm again starts with a local skyline computation at each site. Similar to Algorithm 1, the rest of the tradeoff algorithm still proceeds in rounds. The main difference is that at each round, the parties (sites and the coordinator) first jointly compute $(1/d, 1/(2d))$-quantiles to partition the Euclidean plane to a set of at most d vertical strips, and then instead of computing the (at most 2) points with the global maximum x/y coordinates, the coordinator computes for each non-empty strip the point with the maximum y-coordinate by collecting information from the sites (Line 5-8); after that the parties use these points to compute new skyline points and prune each strip. We call the combination of computing the quantiles and maximum y-coordinates, and finding new skyline points and performing local pruning, one *step* of the computation. The algorithm runs for $(t-1)$ steps, and after that the sites simply send all the remaining points to the coordinator.

Algorithm 2 Communication-Round Tradeoff under Horizontal Partition

Input: S_i is initialized as the point set held by Site i; r (#communication-rounds) is a user-chosen parameter; d is a parameter specified in the analysis.
Output: the global skyline

1: Site i computes its local skyline points and prunes the other (dominated) points in S_i
2: $\ell \leftarrow 0$; $t \leftarrow \lceil r/2 \rceil$
3: **while** $(\ell \leq t-1) \wedge (\exists i \in [s]\ s.t.\ S_i \neq \emptyset)$ **do**
4: All sites and the coordinator jointly compute $(1/d, 1/2d)$-quantiles according to the x-coordinates of points in $\bigcup_{i\in[s]} S_i$ ▷ the quantile points naturally partition the Euclidean plane to d strips
5: **for** each i $s.t.$ $S_i \neq \emptyset$ **do**
6: Site i, for each non-empty strip, sends the point with the largest y-coordinate to the coordinator
7: **end for**
8: The coordinator, for each strip, finds the point with the largest y-coordinate among points received from sites; let Y denote the set of these points among all strips
9: The coordinator computes new skyline points from Y and sends them to each site
10: **for** each i $s.t.$ $S_i \neq \emptyset$ **do**
11: Site i prunes S_i by new global skyline points received from the coordinator
12: **end for**
13: $\ell \leftarrow \ell + 1$
14: **end while**
15: $\forall i \in [s]$, Site i sends S_i to the coordinator
16: The coordinator updates the global skyline using the new points received from sites

We now show the correctness of Algorithm 2 and analyze its costs. The high level intuition on the round efficiency of Algorithm 2 is that at each round, the point with the maximum y-coordinate in each strip will either contribute to the global skyline or help to prune all the points in that strip. Compared with Algorithm 1, one can think that we are trying to prune the whole data set *in parallel*, that is, in each strip of the plane. This will reduce the round complexity at the cost of mildly increasing the total communication cost.

Correctness. The correctness of Algorithm 2 is straightforward: our skyline computation does not prune any point that is not dominated by others. Indeed, up to the $(t-1)$-th step (or, $(2t-2)$-th round), what Algorithm 2 does can be summarized as "sites send candidate global skyline points $\rightarrow$ the coordinator computes new global skyline points from these candidates $\rightarrow$ sites use new skyline points to prune their local datasets". At the $(2t-1)$-th round, sites just send all the remaining unpruned points to the coordinator so that we will not miss any skyline points.

Communication cost. We count the communication cost in two parts. The first part is the communication needed at the first $(t-1)$ steps, and the second part is the total number of remaining points at all sites after the $(t-1)$-th step, which will be sent to the coordinator all at once at the final round.

We first analyze the cost of computing quantiles at each step. We can compute $(\epsilon, \epsilon/2)$-quantiles using the following folklore algorithm: the i-th site (for all $i \in [s]$) sends the coordinator the exact $\epsilon/2$-quantiles Q_i of its local point set S_i. Using $\{Q_1, \ldots, Q_s\}$, the coordinator can answer quantile queries as follows: Given a query rank β, it returns the largest v satisfying

$$\beta - \textstyle\sum_{i\in[s]} rank_i(v) \geq 0,$$

where $rank_i(v) = n_i(v) \cdot (\epsilon/2 \cdot |S_i|)$, and $n_i(v)$ is the number of $\epsilon/2$-quantiles in Q_i that is smaller than v. It is easy to see that $0 \leq \beta - v \leq \epsilon/2 \cdot \sum_{i\in[s]} |S_i|$. The following lemma summarizes the communication cost of this algorithm.

LEMMA 1. *There is an algorithm that computes $(\epsilon, \epsilon/2)$-quantiles in the coordinator model using one round and $O(s/\epsilon)$ communication.*

Thus the communication used for quantile computation can be bounded by $O(sd)$ at each step. The rest communication at each step includes sending local maximums at all strips and new skyline points, which can be bounded by $O(sd)$ we as well. To sum up the total communication in the first part is bounded by $O(sd(t-1))$.

The rest of the analysis is devoted to the second part, that is, to bound the number of the remaining points after the $(t-1)$-th step.

We first assume that the output size k is known, and d will be chosen as a function of k. We will then show how to remove this assumption.

Let $Y_\ell \in [1, d]$ denote the number of new skyline points we find at the ℓ-th step. Observe that in each strip, if the point with the largest y-coordinate is not a skyline point, then the rest of the points in that strip cannot be skyline points and thus are pruned. After the first step, there are at most Y_1 strips having point and each strip has at most $2n/d$ points, so there are at most

$$Y_1 \cdot 2n/d = 2nY_1/d$$

points left. After the second step, there are at most Y_2 strips having point and each strip has at most $2(2nY_1/d)/d$ points, so there are at most

$$Y_2 \cdot 2\left(2nY_1/d\right)/d = 4nY_1Y_2/d^2 \qquad (Y_1 + Y_2 \leq k)$$

points left. After the $(t-1)$-th step, there are at most

$$2^{t-1} n \prod_{\ell\in[t-1]} Y_\ell \,/\, d^{t-1} \qquad \left(\sum_{\ell\in[t-1]} Y_\ell \leq k\right) \quad (1)$$

$$\leq \quad n\left(\frac{2k}{(t-1)d}\right)^{t-1} \quad (2)$$

points left, where from (1) to (2) we have used the AM-GM inequality and the equality holds when all $Y_\ell = k/(t-1)$ ($\ell = 1, \ldots, t-1$). We thus have at most $n(2k/((t-1)d))^{t-1}$ points left at sites after $(t-1)$-th step, and the sites will send all of them in the final (i.e., $(2t-1)$-th) round. Adding two parts together, the total communication cost is bounded by

$$O(sd(t-1)) + n\left(\frac{2k}{(t-1)d}\right)^{t-1}. \quad (3)$$

When

$$d = \frac{2k}{t-1} \cdot \left(\frac{n(t-1)}{2sk}\right)^{1/t}, \quad (4)$$

Expression (3) simplifies to be $O(sk^{(t-1)/t}(n/s)^{1/t})$.

Dealing with unknown k. We now show how to deal with the case that we do not know k at the beginning. A simple idea is to guess k as powers of 2 (i.e., $1, 2, 4, 8, \ldots$), and for each guess, we run our algorithm, and report error if $\sum Y_\ell > k$ at some point, in which case we double the value of k and rerun the algorithm. The correctness of the algorithm still holds. The round complexity, however, may blow up by a factor of $\log k$ in the worst case. We will show that there is a way to preserve the round complexity even when we do not know k at the beginning.

The new idea is to guess k progressively, based on the number of new skyline points found in the previous step. More precisely, we set the guess of k at the ℓ-th step, denoted by k_ℓ, to be $Y_{\ell-1} \cdot (t-1)$, for $\ell \geq 2$; and we set $k_1 = t-1$ to begin with. Now at the ℓ-th step we use

$$d = d_\ell = \frac{2k_\ell}{t-1} \cdot \left(\frac{n(t-1)}{2s}\right)^{1/t} \tag{5}$$

strips for the pruning. Note that (5) is very similar to (4), where we have replaced the first k in (4) by k_ℓ and removed the second k in (4).

Similar to (2), after $(t-1)$-th step, there are at most

$$2^{t-1} n \prod_{\ell \in [t-1]} (Y_\ell / d_\ell)$$

points left, and consequently the total communication cost is bounded by

$$s \sum_{\ell \in [t-1]} d_\ell + 2^{t-1} n \prod_{\ell \in [t-1]} (Y_\ell / d_\ell). \tag{6}$$

We now bound the two terms in (6) separately. We first have

$$\begin{aligned} s \sum_{\ell \in [t-1]} d_\ell &= \frac{2s \sum_{\ell \in [t-1]} k_\ell}{(t-1)} \cdot \left(\frac{n(t-1)}{2s}\right)^{1/t} \\ &\leq 2sk \cdot \left(\frac{n(t-1)}{2s}\right)^{1/t}. \end{aligned} \tag{7}$$

where we have used the inequality

$$\sum_{\ell \in [t-1]} k_\ell = (t-1)\left(1 + \sum_{\ell \in [t-2]} Y_\ell\right) \leq (t-1)k.$$

For the second term, we have

$$\begin{aligned} 2^{t-1} n \prod_{\ell \in [t-1]} (Y_\ell / d_\ell) &= n \frac{Y_1 Y_2 \cdots Y_{t-1}}{Y_1 Y_2 \cdots Y_{t-2}} \left(\frac{n(t-1)}{2s}\right)^{t/(t-1)} \\ &= n Y_{t-1} \cdot \frac{2s}{n(t-1)} \left(\frac{n(t-1)}{2s}\right)^{1/t} \\ &\leq \frac{2sk}{(t-1)} \left(\frac{n(t-1)}{2s}\right)^{1/t}. \end{aligned} \tag{8}$$

By (7) and (8), the total cost is bounded by $O(sk(n/s)^{1/t}) = O(sk(n/s)^{1/\lceil r/2 \rceil})$, as claimed.

Running time. The time cost at each site involves three parts: the computation of the local skyline, the computation of local quantiles, and the point prunings. Computing local skyline again cost $O(n_i \log n_i)$ where n_i is the number of points at the i-th site. The cost of point prunings can again be made linear in n_i for the same reason as that in Algorithm 1. Now we analysis the time cost of computing the local quantiles. Since points are sorted after the local skyline computation, computing (exact) local quantiles needs $O(d_\ell)$ time at the ℓ-th step. Thus the total time is bounded by $\sum_{\ell \in [t-1]} d_\ell = O(k(n/s)^{1/t}) = O(k(n/s)^{1/\lceil r/2 \rceil})$.

The running time at the coordinator also consists of three parts: the computation of approximate global quantiles at each step, the computation of new skyline points from the first step to the $(t-1)$-th step, and the computation of skyline points (output) at the end. The observation is that at each step, for each of the three tasks, the running at the coordinator can be asymptotically bounded by the number of points it receives from all sites in that step, and thus the total running time at the coordinator is asymptotically upper bounded by the total communication cost.

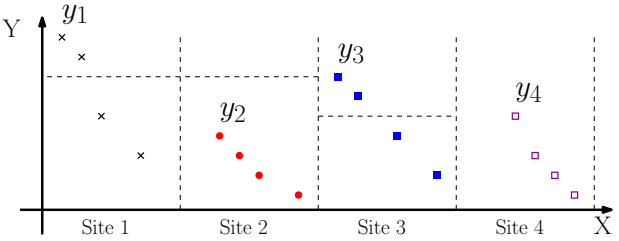

Figure 3: Algorithm for Sorted Dataset

2.1.3 An Algorithm for Sorted Datasets

In this section we show that we can do much better if the data points are partitioned to the s sites in a *sorted* order with respect to the x-coordinate (or, the y-coordinate). In other words, let $x_1 \leq x_2 \leq \ldots \leq x_{s-1}$ be $(s-1)$ split points, and let $x_0 = -\infty$ and $x_s = \infty$. The i-th site gets all points between $(x_{i-1}, x_i]$. See Figure 3 for an illustration. As mentioned in the introduction, in the coordinator model we cannot afford to repartition the dataset. Therefore the algorithm presented in this section is only useful for settings where the data has already been sorted among the sites.

THEOREM 3. *For n points in the 2-dimensional Euclidean space partitioned among the s sites in the coordinator model in the sorted manner according to their x-coordinates or y-coordinates, there exists an algorithm for computing the skyline that uses $O(k+s)$ communication and 2 rounds, where k is the output size, that is, the number of points in the skyline. The total running time at the i-th site is $O(n_i \log n_i)$ where n_i is the number of points at the i-th site, and that at the coordinator is $O(k+s)$.*

We described our algorithm for sorted datasets in Algorithm 3. Let us explain in words. Same as before, each site first does a local pruning and computes its local skyline. In the first round, the i-th site sends the coordinator the point which has the largest y-coordinate, denoted by y_i. In the second round, the coordinator for each $i \in [s]$ computes the value z_i which is the maximum value among $\{y_{i+1}, \ldots, y_s\}$, and sends z_i to the i-th site. The i-th site then prunes all its local points with y-coordinate smaller than z_i, and sends the rest of its points to the coordinator.

We now show the correctness of this algorithm and analyze its costs. The claim is that the points in S_i with y-coordinate

Algorithm 3 Sorted Datasets under Horizontal Partition

Input: S_i is initialized as the point set held by Site i
Output: the global skyline

1: Site i computes its local skyline points and discards the other points in S_i
2: Site i sends the point with the largest y-coordinate, denoted by y_i, to the coordinator
3: The coordinator computes for $\forall i \in [s], z_i = \max\{y_{i+1}, \ldots, y_s\}$, and sends z_i to Site i
4: Site i prunes S_i using z_i, and sends the rest points to the coordinator
5: The coordinator updates the global skyline using the new points received from sites

larger than z_i, denoted by P_i, must on the global skyline. Indeed, points in P_i can not be dominated by any point in $S_1, \ldots, S_{i-1}$ since all points in S_i have x-coordinates larger than those points in $S_1, \ldots, S_{i-1}$. On the other hand, points in P_i cannot be dominated by any point in $S_{i+1}, ..., S_s$ since all points in P_i have y-coordinates larger than the those points in $S_{i+1}, \ldots, S_s$. The communication of the algorithm includes sending y_i and z_i (costs $2s$) plus sending skyline points (costs k). The running time at each site is dominated by the local skyline computation, and the time cost at the coordinator is clearly $O(k+s)$ ($O(s)$ for the first round and $O(k)$ for the second round).

2.2 Lower Bounds

2.2.1 *Infinite Rounds*

We prove a lower bound for the infinite-round case by a reduction from a communication problem called s-DISJ. The lower bound matches the upper bound by Algorithm 1 up to a logarithmic factor which counts the number of bits used to represent a point in the Euclidean plane.

In s-DISJ, each of the s sites gets an m-bit vector. Let $X_i = (X_i^1, \ldots, X_i^m)$ be the vector the the i-th site gets. We can view the whole input as an $s \times m$ matrix X with X_i ($i \in [s]$) as rows. The s-DISJ problem is defined as follows:

$$s\text{-DISJ}(X_1, \ldots, X_s) = \begin{cases} 1, & \text{if there exists a } j \in [m] \text{ s.t.} \\ & \forall i \in [s], X_i^j = 1, \\ 0, & \text{otherwise.} \end{cases}$$

LEMMA 2 ([5]). *Any randomized algorithm for s-DISJ that succeeds with probability 0.51 has communication cost $\Omega(sm)$. The lower bound holds even when we allow an infinite number of communication rounds.*

The Reduction. Given the m-bit vector X_i for s-DISJ, the i-th site first converts it to a $2m$-bit vector X_i' as follows: each 0 bit will be converted to 01, and each 1 bit will be converted to 10. For example, when $m = 5$ and $X_i = 10101$, X_i' should be 1001100110. The next step is to convert X_i' to a staircase. This step is illustrated in Figure 4. We can "embed" the staircase into an $m \times m$ grid. The staircase starts from the top-left point of the grid, and grows in $2m$ steps. In the ℓ-th step, if the ℓ-th coordinate of X_i' is 0, then the staircase grows one step horizontally rightwards; otherwise if the ℓ-th coordinate is 1, then the staircase grows one step vertically downwards.

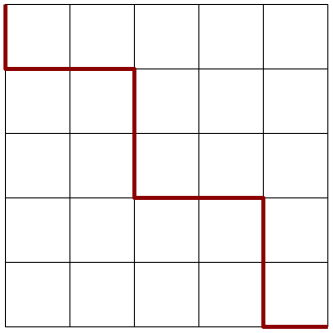

Figure 4: Translating vector $X_1' = 1001100110$ ($m = 5$) to a staircase.

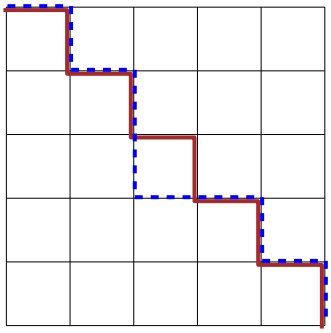

Figure 5: The solid red skyline corresponds to the case that s-DISJ$(X_1, \ldots, X_s) = 0$, and the dash blue skyline corresponds to the case that s-DISJ$(X_1, \ldots, X_s) = 1$.

The observation is that if we create s staircases using $X_1', \ldots, X_s'$, then the skyline of the union of these s staircases is closely related to the value of s-DISJ$(X_1, \ldots, X_s)$: If s-DISJ$(X_1, \ldots, X_s) = 0$, then the skyline will be in the form of the red curve in Figure 5; otherwise, the skyline will be different from the red curve (e.g., be the blue curve in Figure 5 if the 3rd coordinates of $X_1, \ldots, X_s$ are all 1). This is because for each column $j \in [m]$ in the grid, as long as there is one $i \in [s]$ such that the j-th coordinate of X_i is 0, or the $(2j-1)$-th and $(2j)$-th coordinates of X_i' is 01, the skyline within the j-th column of the grid will be like "$\urcorner$"; otherwise if for all $i \in [s]$ the j-th coordinate of X_i is 0, then the skyline within the j-th column of the grid will be like "$\llcorner$". The other direction also holds, that is, if the skyline is in the form of the red curve, then s-DISJ$(X_1, \ldots, X_s) = 0$; otherwise, s-DISJ$(X_1, \ldots, X_s) = 1$.

When the size of the skyline is k, we set $m = k$ according to our reduction, and obtain the following theorem.

THEOREM 4. *Any randomized protocol for computing skyline in the coordinator model with horizontal partition that succeeds with probability 0.51 has communication cost $\Omega(sk)$ bits, where k is number of points in the skyline. The lower bound holds even when we allow an infinite number of communication rounds.*

2.2.2 *One Round*

We now prove an $\Omega(n)$ communication lower bound for the case that the algorithm needs to finish in one round. This proof is simpler than the infinite-round case – we only need

two sites to participate in the game. We assign Site 1 an m-bit vector u, which can be converted to a staircase in the $m \times m$ grid just like the infinite-round case, and assign Site 2 one bit v, which is translated to the upper-right conner point of the grid if $v = 1$ and the lower-left conner point of the grid if $v = 0$. We have the following simple observation, which holds because if we change one bit of the vector m from 0 to 1, we will change a "$\urcorner$" to "$\llcorner$" in the staircase, and thus change the skyline; same for replacing a bit 1 to 0.

OBSERVATION 1. *If the coordinator needs to compute the global skyline, and $v = 0$, then it needs to learn the vector u exactly.*

We immediately have the following lemma.

LEMMA 3. *Any one round randomized algorithm for the coordinator to learn u exactly with probability 0.51 has communication cost $\Omega(m)$.*

Note that when $v = 1$, the skyline only consists of a single point in the upper right conner. Therefore the output size k is not directly related to the value m; we thus can set $m = \Omega(n)$.

THEOREM 5. *Any one round randomized algorithm for computing skyline in the coordinator model with horizontal partition that succeeds with probability 0.51 has communication cost $\Omega(n)$ bits, where n is the total number of points held by all sites.*

3. VERTICAL PARTITION

In the vertical partition, each site holds a single coordinate of all the data points. Since we consider points in the 2-dimensional Euclidean space in this paper, we only have two sites, named Alice and Bob. We store the ID (key) of each point on both sites. More precisely, Alice has a set of tuples $\{(x_{p_1}, \mathsf{ID}_{p_1}), (x_{p_2}, \mathsf{ID}_{p_2}), \ldots\}$, where x_{p_i} is the x-coordinate of point p_i and ID_{p_i} is the ID of p_i; and Bob has a set of tuples $\{(\mathsf{ID}_{q_1}, y_{q_1}), (\mathsf{ID}_{q_2}, y_{q_2}), \ldots\}$, where y_{q_i} is the y-coordinate of point q_i. We can join the two tables to recover the information of all points.

3.1 A Strong Lower Bound

As mentioned in the introduction, in the vertical partition case there is a strong communication lower bound, which states that in the worst case sites have to send almost everything to the coordinator. We prove the lower bound by a reduction from a classical problem in communication complexity called two-party set-disjointness (denoted by 2-DISJ). In the 2-DISJ, we have two parties Alice and Bob. Alice holds a set of numbers $A \subseteq [n]$, and Bob holds a set of numbers $B \subseteq [n]$, and they want to jointly compute

$$\text{2-DISJ}(A, B) = \begin{cases} 1, & A \cap B \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

FACT 1 ([3]). *Any randomized protocol for 2-DISJ that succeeds with probability 0.51 has communication cost $\Omega(n)$.*

The Reduction. Alice uses her input A to create the x-coordinates of the n points as follows: for each $a \in A$, she creates a point $(2, a)$ where 2 is the x-coordinate and a is the ID of the point; and for each $a \in [n] \backslash A$, she creates a point $(1, a)$. Bob uses his input B to create the y-coordinates of the n points in a similar way: for each $b \in B$, he creates a point $(b, 2)$ where 2 is the y-coordinate and b is the ID of the point; and for each $b \in [n] \backslash B$, he creates a point $(b, 1)$.

It is easy to see that if $A \cap B \neq \emptyset$, then there is at least one point in the $(2, 2)$ position, which will be the only point in the skyline; otherwise, the skyline will consist of two points $(1, 2)$ and $(2, 1)$. By Fact 1 we have

THEOREM 6. *Any randomized algorithm for computing skyline in the distributed model with vertical partition that succeeds with probability 0.51 has communication cost $\Omega(n)$, regardless how many rounds the algorithm uses.*

3.2 A Simple Heuristic

Given the strong lower bound result in Section 3.1, we propose in this section a heuristic which is communication efficient on real-world datasets. Our algorithm can be thought as a batched pruning version of the Threshold-Algorithm based approach [2, 12], but has been tailored for bounded-round distributed computation. We call our heuristic *interactive pruning*; the pseudocode is described in Algorithm 4.

We now explain our algorithm in words. Let r be the round budget and ρ be the parameter denoting the number of groups we use for the algorithm. We say a point p is *recovered* at the coordinator if both the x and y coordinates of p are known to the coordinator.

Alice and Bob first sort decreasingly all points according to the x-coordinates and y-coordinates respectively, and then partition the points to ρ groups, denoted by $G_{x1}, \ldots, G_{x\rho}$ and $G_{y1}, \ldots, G_{y\rho}$ respectively. The coordinator tries to learn both coordinates of the points that possibly lie in the skyline progressively. It maintains four variables f_x, f_y, l_x, l_y: f_x and f_y stand for the indices of the next groups Alice and Bob will send to the coordinator respectively; and groups with indices at least l_x and l_y at Alice and Bob respectively have already been pruned. Observe that if $f_x \geq l_x$ or $f_y \geq l_y$, then the coordinator must have recovered all the skyline points, and thus the algorithm can terminate. Based on this fact, the coordinator always chooses the site with the smaller *gap value* $l_x - f_x$ or $l_y - f_y$, and asks for the information of its next group G_{xf_x} or G_{yf_y}.

Algorithm 4 proceeds in three stages. The first stage consists of two rounds (Line 3 to 5). The goal of the first stage is to recover the first groups G_{x1} and G_{y1} at Alice and Bob, and update the variables l_x and l_y to determine the next group to recover. The second stage consists of $\lfloor (r-4)/2 \rfloor$ steps each of which consisting of two rounds (Line 6 to 10). The goal of the second stage is to recover points progressively from the site which has a smaller gap value. The algorithm will stop early if at some step all points have been either pruned or recovered at a site. The goal of the third stage is to recover all groups that are still *not* recovered or pruned after the first two stages. This is done at the coordinator by first asking the site with the smaller gap value to get the IDs of points in these groups together with one of the two coordinates, and then contacting the other site for the other coordinates of these points. We include a running example for Algorithm 4 in Appendix A.

The correctness of the algorithm is straightforward since when our algorithm stops, all points at one site are either recovered or pruned. The local running time at each site is dominated by the initial sorting which is $O(n \log n)$, and

Algorithm 4 Interactive Pruning under Vertical Partition

Input: S_1 is the set of the x-coordinates of points held by Alice, and S_2 is the set of the y-coordinates of points held by Bob. ρ (number of groups) and r (round budget) are two user-chosen parameters.

Output: the global skyline

1: Alice sorts S_1 decreasingly and then partitions it to ρ groups of equal size, denoted by $\mathcal{G}_x = \{G_{x1}, \ldots, G_{x\rho}\}$. Bob does the same thing on S_2 and gets groups $\mathcal{G}_y = \{G_{y1}, \ldots, G_{y\rho}\}$. Alice and Bob send the splitting coordinates of the group partitions to the coordinator. Given a point p, let $g_x(p)$ be the index of the group in $\mathcal{G}_x$ that contains p, and $g_y(p)$ be the index of the group in $\mathcal{G}_y$ that contains p
2: The coordinator maintains a set of points R denoting the recovered points, and four variables: f_x, f_y, l_x, l_y.
3: The coordinator communicates with Alice and Bob to learn both the x-coordinates and y-coordinates of points in G_{y1} and G_{x1}, adds them to R, and then updates $l_x \leftarrow \min\{g_x(p) \mid p \in G_{y1}\}+1$ and $l_y \leftarrow \min\{g_y(p) \mid p \in G_{x1}\}+1$
4: The coordinator sets $f_x \leftarrow 2$, $f_y \leftarrow 2$
5: $j \leftarrow 1$ ▷ index of steps; one step uses two rounds
6: **while** $(2j \leq r-4) \wedge (l_x > f_x) \wedge (l_y > f_y)$ **do**
7: If $l_x - f_x \leq l_y - f_y$, the coordinator requests group G_{xf_x} from Alice and updates $f_x \leftarrow f_x+1$, else it requests G_{yf_y} from Bob and updates $f_y \leftarrow f_y+1$
8: The coordinator recovers and adds points in G_{xf_x} (or G_{yf_y}) to R, and updates $l_y \leftarrow \min\{l_y, \min\{g_y(p) \mid p \in G_{xf_x}\}+1\}$ (or $l_x \leftarrow \min\{l_x, \min\{g_x(p) \mid p \in G_{yf_y}\}+1\}$)
9: $j \leftarrow j+1$
10: **end while**
11: **if** $(l_x > f_x) \wedge (l_y > f_y)$ **then**
12: The coordinator compares $l_x - f_x$ and $l_y - f_y$; w.l.o.g., assume $l_x - f_x$ is smaller. The coordinator sends f_x and l_x to Alice. Alice then sends the x-coordinates and IDs of points (denoted by P) in $G_{xf_x}, \ldots, G_{x(l_x-1)}$ to the coordinator.
13: The coordinator sends the IDs of points in P to Bob. Bob then sends the y-coordinates of points in P to the coordinator. The coordinator adds points in P to R
14: **end if**
15: The coordinator computes and outputs the skyline of R

the time cost at the coordinator is dominated by the final skyline computation which is also $O(n \log n)$.

We comment that ρ is a parameter that we need to choose at the beginning of the algorithm; we will discuss how to choose ρ in the experiments in Section 4.

A comparison to previous algorithms. Our algorithm, the `BDS` algorithm and its improvement the `IDS` algorithm proposed in [2], and the `PDS` algorithm proposed [12] are all based on the Threshold-Algorithm (TA) [7]. A clear difference is that `BDS`, `IDS` and `PDS` recover one point in each step at the coordinator, while we recover at least one group of points in each step which helps to significantly reduce the round cost. The second major difference is that to improve the basic TA-base algorithm `BDS`, both `IDS` and `PDS` share the idea of estimating the *most probable terminating point*, and then using it to decide which site to access next. `IDS` finds the most probable terminating point by calculating the score of each point which is the sum of the differences between each of its coordinates and the coordinate of the last point recovered by the sorted access at the respective site. `PDS` uses linear regression to estimate the rank of each point, and the point with the lowest rank is considered to be the most probable terminating point. In our algorithm, instead of looking for the most probable terminating point, we focus on the number of unrecovered and unpruned groups remaining at each site. We choose to request a new group of points from the site with the smaller number of remaining groups, with the purpose of terminating the algorithm earlier.

It is well-known that the Threshold-Algorithm [7] is *instance optimal* in terms of number of points probed, and thus if we only measure the number of points recovered at the coordinator in the process, our batched pruning approach has no advantage compared with the individual point prunings in `BDS`, `IDS` and `PDS`. However, as we shall see in the experiments (Section 4), batched pruning can not only significantly reduce the round cost, but also bring down the overall communication cost in some cases. This is because in each step the coordinator needs to send a message to request points from the sites in the sorted access, and consequently the round cost will influence the communication cost as well.

4. EXPERIMENTS

In this section we present the experimental studies of our proposed algorithms. We have implemented all algorithms in C++. All experiments were conducted on a laptop with Inter Core i7 running Windows 7 with 4GB memory.

The Datasets. We use both synthetic and real-world datasets. We generated three synthetic datasets following the standard literature [4]. The data partition among the sites will be described in Section 4.1 (horizontal partition) and Section 4.2 (vertical partition) respectively.

- `INDI`(independent): We generate 20 million points. For each point we generate each of its coordinates independently uniformly at random from $[0, 1]$
- `CORR`(correlated): We select 200 thousand lines (denoted by L_1) perpendicular to the line from $(0, 0)$ to $(1, 1)$, where the intersections follow the standard normal distribution. Next, for each line in L_1 we pick 100 points also following the standard normal distribution. We have 20 million points in total.
- `ANTI`(anti-correlated): We select 200 thousand lines (denoted by L_2) perpendicular to the line from $(0, 1)$ to $(1, 0)$, where the intersections follow the standard normal distribution. Next, for each line in L_2 we pick 100 points also following the standard normal distribution. We have 20 million points in total.

We make use of the following real-world datasets.

- `Airline`: this dataset contains 3 million the airline itineraries and prices between 30 U.S. major cities in 2015 first quarter.[2] We choose (minus) *fare* and (minus) *fare-per-mile* as the two attributes in x and y coordinates. The skyline points are considered to be economic flights.

[2]Available at `http://www.transtats.bts.gov`.

- `Household` [11]: this dataset contains 2 million household electric power consumption records gathered between December 2006 and November 2010. We choose *voltage* and *intensity* as the two attributes in x and y coordinates. The skyline points represent those households that are recommended to pay attention to the energy efficiency.
- `Covertype` [11]: this dataset contains 500 thousand natural statistics from four wilderness areas located in the Roosevelt National Forest of northern Colorado. We choose *elevation* and *slope* as the two attributes in x and y coordinates. The skyline points represent areas that may have interesting geological behaviors.

4.1 Horizontal Partition

Algorithms. We compare the following algorithms in the case of horizontal partition.

- `Naive`: the single round algorithm in which each site computes and sends its local skyline to the coordinator for a merge.
- `Optimal`: Algorithm 1 in Section 2.1.1. The algorithm that achieves the optimal communication cost.
- `Tradeoff`: Algorithm 2 in Section 2.1.2. The algorithm that gives a smooth communication-round tradeoff.
- `AGiDS`: We use the `AGiDS` algorithm proposed in [15] as a comparison. To make `AGiDS` fit in our model, we use the following version of the original algorithm: At the beginning, the coordinator and sites share the information of a grid in which each cell represents a range in the x/y-axes (equal width partition). In the first round (the *planing phase*), sites send the information of non-empty cells to the coordinator. In the second round (the *execution phase*), the coordinator finds the cells that may contribute to global skyline and sends the information to sites, and then sites send points in these cells to the coordinator. Finally the coordinator computes the skyline of received points as the output. We comment that this algorithm is also very similar to the *relaxed skyline* algorithm proposed in [1]. We choose the grid to be 20×20 (total 400 cells) in our experiments which works the best under our settings.
- `FDS`: We use the `FDS` algorithm proposed in [20] as a comparison. The original algorithm proceeds in *iterations*. To make it fit in our model, we use three rounds for each iteration: In the first round (the *voluntary phase*), each site sends the top κ points with the largest scores $(x + y)$ to the coordinator. In the second round (the *compulsory and computation phase*), the coordinator calculates the minimum score (denoted by $F_{\min}$) from received points and sends it to each site. Each site then sends all its local points that have larger scores than $F_{\min}$ to the coordinator. The coordinator updates the global skyline with points received in the first two rounds. In the third round (the *feedback phase*), the coordinator calculates and sends each site a feedback, which consists of points that are guaranteed to dominate at least ℓ points in that site. And then each site does a local pruning. κ and ℓ are two parameters in `FDS`; we choose the optimal values $\kappa = 1$ and $\ell = 1$ as reported in [20].

Data Partition. For the three synthetic datasets, we partition points randomly to 20 sites. For `Airline`, we partition the data records in the same city to the same site; we thus have 30 sites. For `Household`, we partition the data collected in every two consecutive months to the same site; 20 sites in total. For `Covertype`, we just randomly partition the data records to 20 sites.

Results and Discussions. Figure 6 and Figure 7 show the communication and round costs of the five tested algorithms on the three synthetic datasets and three real-world datasets mentioned above.

We observed that by using three rounds, the communication cost of `Tradeoff` is 25%-44% of that of `Naive` on synthetic datasets and 22%-31% on real-world datasets, and is also close to `Optimal` which uses more rounds. We noticed that the advantage of `Tradeoff` against `Naive` is larger in real-world datasets (even of smaller sizes) than synthetic ones. This is because it is more likely in real-world datasets that a point in one site dominates most of the local skyline points in another site. We observe that in `Household`, by using 5 rounds `Tradeoff` needs even less communication than `Optimal`. This is possible since `Optimal` is just asymptotically optimal in the *worst case*.

The communication cost of `AGiDS`, which uses two rounds, is similar to `Naive` on `Airline` and `Household` and even worse on `INDI`, `CORR`, and `Covertype`, and is consequently much worse than `Tradeoff` and `Optimal`. `FDS` uses the largest amount of rounds (even more than `Optimal`), but its communication is worse than both `Tradeoff` and `Optimal` on all synthetic and real-world datasets.

Figure 8 and Figure 9 show the running time of the tested algorithms (excluding the cost the common local skyline computation at the beginning) on synthetic and real-world datasets respectively.[3] Generally speaking the running time of all algorithms are similar. On the `ANTI` dataset, the running time of `Optimal` and `FDS` are clearly worse than others. This is because they need many more rounds than other algorithms on `ANTI`.

Summary. `Tradeoff` achieves noticeable communication cost reductions than `AGiDS` and `Naive` by using one or two more rounds; its performance is very close to the theoretical optimal algorithm `Optimal` in the communication cost but is much more efficient in rounds. On the other hand, `AGiDS` does not have an advantage against `Naive` in communication but uses one more round, and the performance of `FDS` is clearly dominated by `Tradeoff` and `Optimal`. All algorithms have similar time costs since the computation of the local skylines dominates the other costs.

4.2 Vertical Partition

We compare the following algorithms in the vertical partition case.

- `Prune`: Algorithm 4 in Section 3.2.
- `BDS` and `IDS`: We use the `BDS` and `IDS` algorithms proposed in [2] as comparisons. See Section 3.2 for a brief description of these two algorithms.

[3]The local skyline computation takes about 6 seconds on all the three synthetic datasets and 1.1, 0.6, 0.2 seconds on the `Airline`, `Household` and `Covertype` respectively. This cost in fact dominates the other time costs.

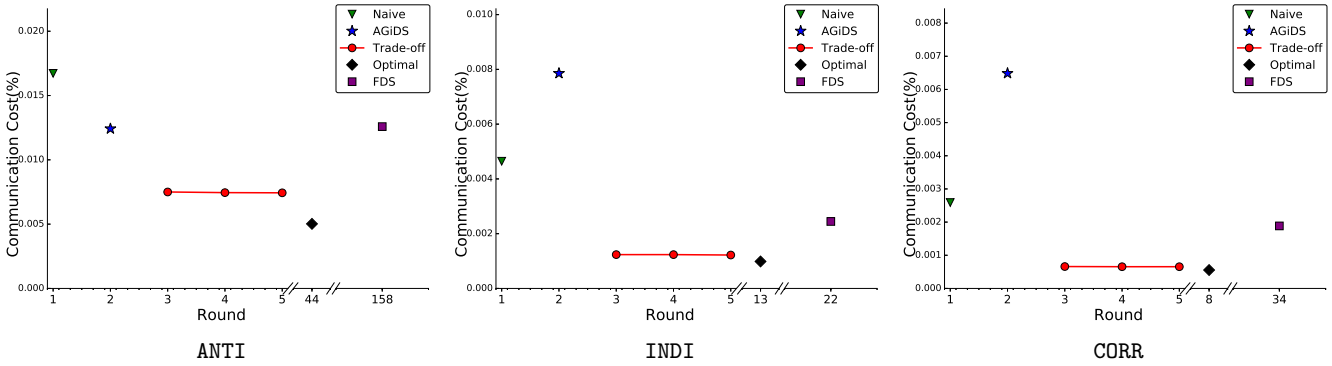

Figure 6: Communication and round costs on synthetic datasets under horizontal partition

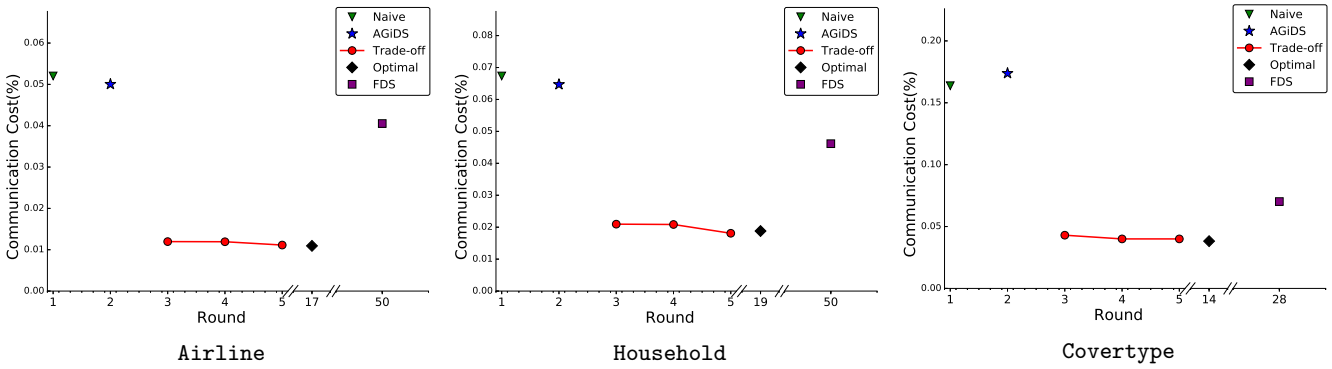

Figure 7: Communication and round costs on real-world datasets under horizontal partition

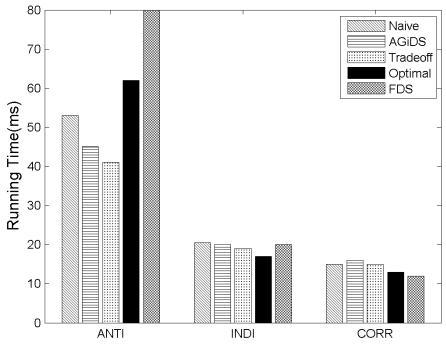

Figure 8: Running time on synthetic datasets under horizontal partition (in `Tradeoff` we use 5 rounds). The time for local skyline computation at the beginning is not included.

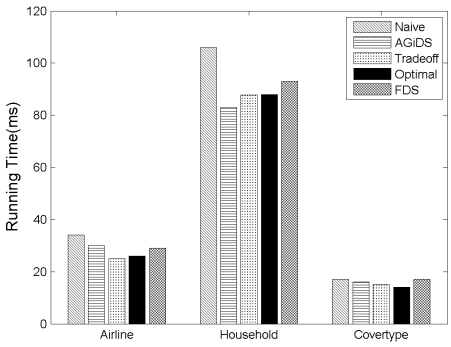

Figure 9: Running time on real-world datasets under horizontal partition (in `Tradeoff` we use 5 rounds). The time for local skyline computation at the beginning is not included.

- `PDS`: We use the `PDS` algorithm proposed in [12] as a comparison. See Section 3.2 for a brief description of this algorithm.

Figure 10 and Figure 11 show the communication and round costs of the four tested algorithms on the three synthetic datasets and three real-world datasets mentioned above. We observed that by using at most eight rounds, the communication cost of `Prune` is similar or even smaller than that of `BDS`, `IDS` and `PDS`, which use significantly more rounds. It is surprising to see that `PDS` is even worse then `BDS` in most cases; the original paper [12] did not make this comparison on points in the 2D Euclidean space. One reason that `Prune` can achieve better communication costs is that `IDS` and `PDS` have significantly larger round complexities, and at every two rounds the coordinator needs to send a message to a site to request for the next point, which can be considered as an overhead of the algorithms in the distributed setting. We also observe that `Prune` recovers similar number of points as `BDS`, `IDS`, and `PDS`; see Table 1 for the details.

In Figure 10 and Figure 11 we have set ρ to be 1000, 10000, 100000 for three synthetic datasets `ANTI`, `INDI` and `CORR`, and 500 for all three real-world datasets. These values are not

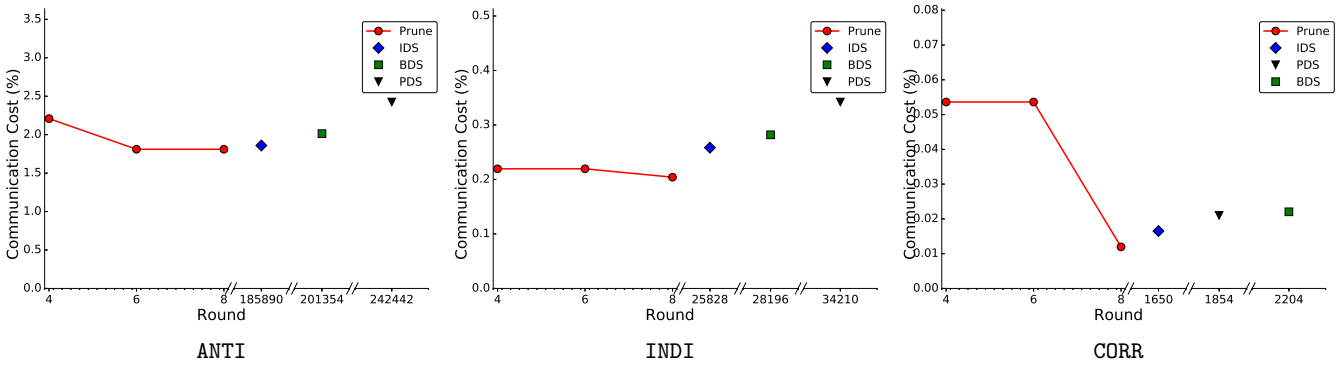

Figure 10: Communication and round costs on synthetic datasets under vertical partition; for `Prune` we use parameter $\rho = 1000, 10000, 100000$ for `ANTI`, `INDI`, `CORR` datasets respectively)

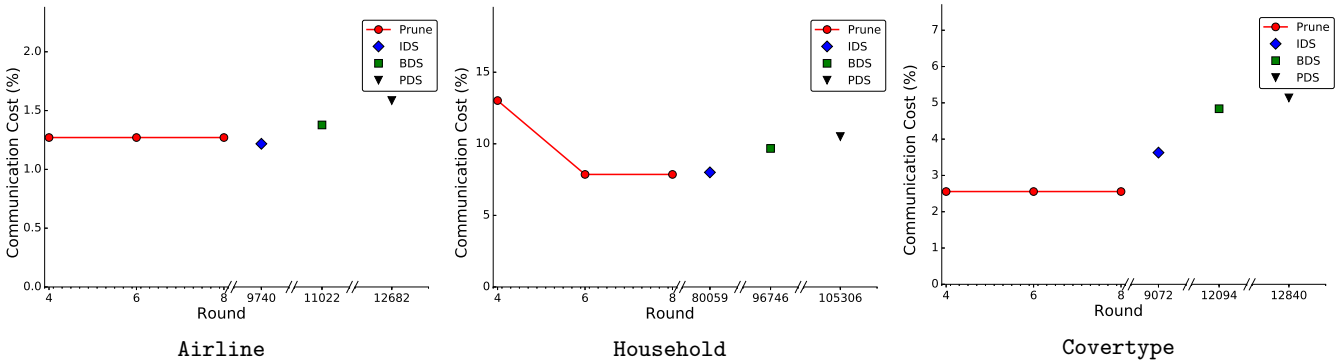

Figure 11: Communication and round costs on real-world datasets under vertical partition; for `Prune` we use parameter $\rho = 500$ for all `Airline`, `Household`, `Covertype` datasets)

	Prune	IDS	BDS	PDS
ANTI	136094	92945	100677	121221
INDI	13609	12914	14098	17105
CORR	855	825	1102	927
Airline	6776	4870	5511	6341
Household	52407	40029	48373	52653
Covertype	5258	4536	6047	6420

Table 1: Number of recovered points at the coordinator. In `Prune`, ρ is chosen to be 1000, 10000, 100000 for `ANTI`, `INDI` and `CORR` respectively, and 500 for `Airline`, `Household` and `Covertype`; the round budget r is set to be 8.

chosen optimally for each dataset, but follow some general rules we have concluded from another set of experiments, which we discuss next.

Figure 12 and Figure 13 show how the communication cost changes when varying ρ. We observe that for a fixed round budget, the communication cost curves generally obey a "U" shape: the cost is very large when ρ is very small; this is because the group size is too large so many recovered points are wasted in the first few rounds. The cost then deceases dramatically as ρ increases, becomes stable for a while, and then increases again. We also observe that the stable range of ρ becomes larger when the round budget r becomes larger, which makes the choice of a good ρ more flexible. While the best ρ inherently depends on the distribution of the dataset, the lesson we have learned from Figure 12 and Figure 13 is that ρ should be small (e.g., 1000) if the data distribution is close to `ANTI`, large (i.e., 100000) if the distribution is close to `CORR`, and medium (i.e., 10000) if the distribution is close to `INDI`. It is reasonable to believe that most real-world datasets are on the `ANTI` end, since an item which is strong on one attribute may be weak on the other (think about the price and location of hotels). Our experiments on three real-world datasets `Airline`, `Household`and `Covertype` show that $\rho = 500$ can be a good default choice in practice.

Figure 14 and Figure 15 show the time costs of tested algorithms on synthetic and real-world datasets respectively. We do not take into account the initial sorting time which is common to all algorithms.[4] Generally speaking, the running times of all algorithms are similar. This is because all of the three algorithms are based on the TA approach. `PDS` has the worst time performance in most cases because it needs to recover the largest number of points *and* uses the most number of rounds.

Summary. To conclude, `Prune` achieves a significant reduction on the round costs (3-4 orders of magnitude) while preserves communication costs compared with `IDS`, which is the best among the three existing algorithms. The running times of all the tested algorithms are similar.

[4]The sorting time is about 5 seconds on all the three synthetic datasets, and 0.9, 0.5, 0.15 seconds on `Airline`, `Household` and `Covertype`.

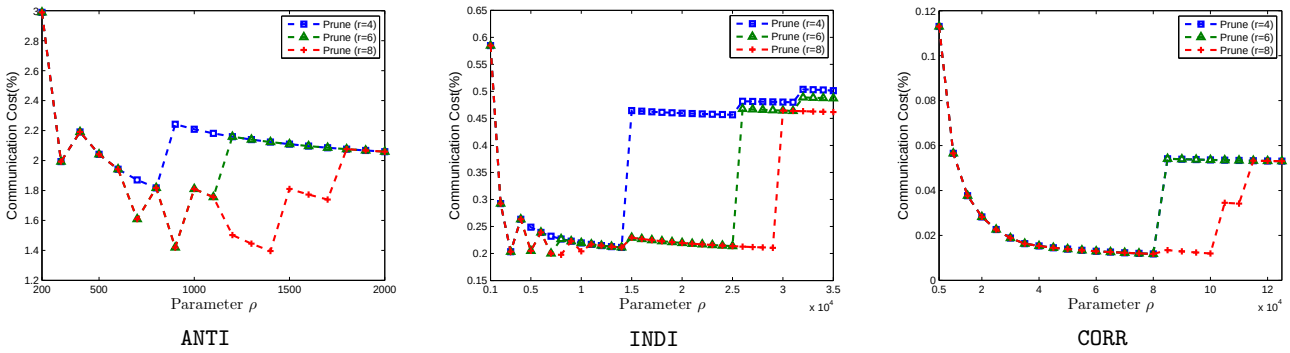

Figure 12: Communication cost versus parameter ρ on synthetic dataset under vertical partition

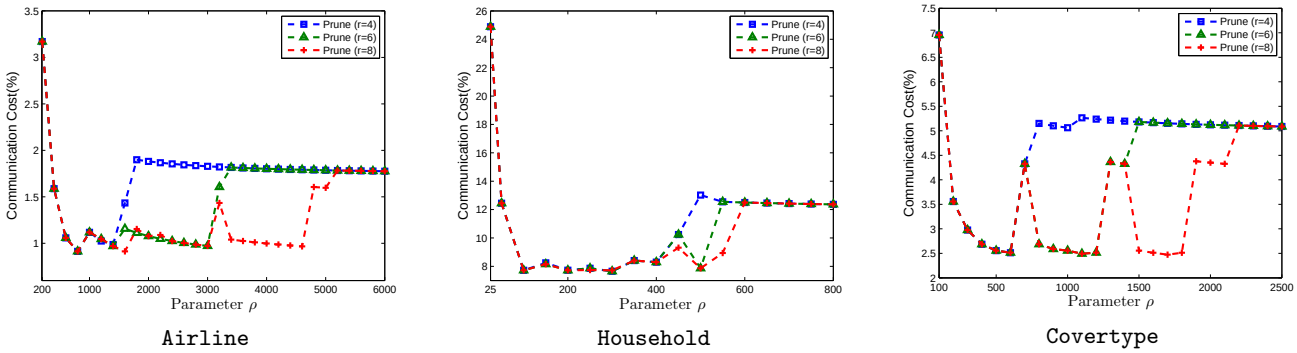

Figure 13: Communication cost versus parameter ρ on synthetic dataset under vertical partition

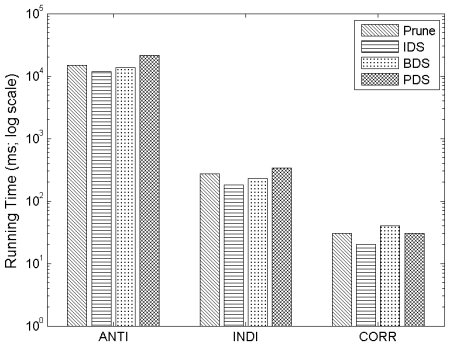

Figure 14: Running time on synthetic datasets under vertical partition; for `Prune` we use parameter $\rho = 1000, 10000, 100000$ for `ANTI`, `INDI`, `CORR` datasets respectively, and $r = 8$ rounds). The time for local skyline computation at the beginning is not included.

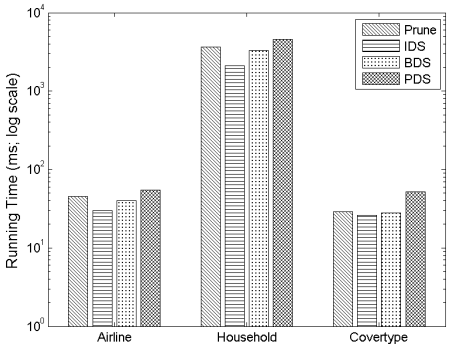

Figure 15: Running time on real-world datasets under vertical partition; for `Prune` we use parameter $\rho = 500$ for all `Airline`, `Household`, `Covertype` datasets, and $r = 8$ rounds). The time for local skyline computation at the beginning is not included.

5. CONCLUSION

In this paper we propose a set of algorithms for computing skylines on distributed data. We consider both horizontal and vertical data partitions. For horizontal partition we give an algorithm that achieves the optimal communication cost. We also propose an algorithm with a smooth communication-round tradeoff. We show experimentally that our algorithms (with theoretical guarantees) significantly outperform existing heuristics in the communication cost and/or round cost. For vertical partition, we first show a strong lower bound, and then give a practical heuristic. Our heuristic significantly outperforms existing ones in the round cost, while achieves similar or better communication cost. Further work includes designing algorithms for points in the higher dimensional Euclidean spaces for both horizontal and vertical data partitions.

6. REFERENCES

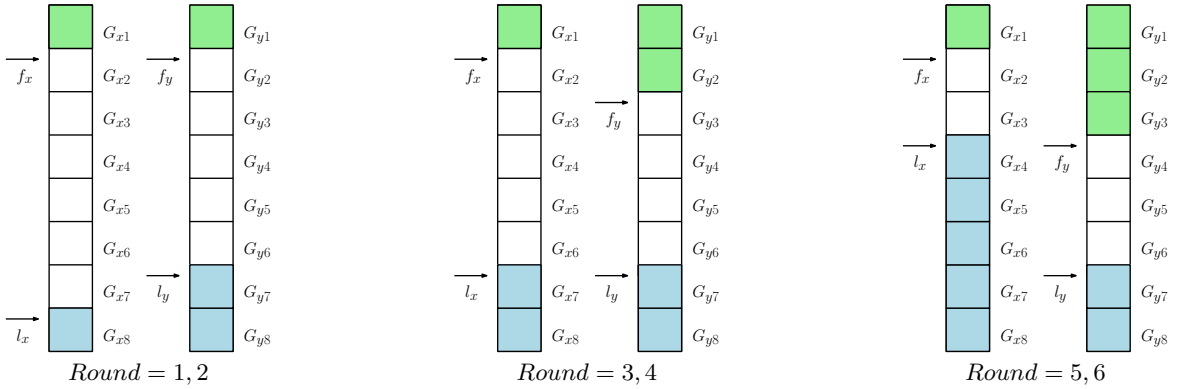

Figure 16: A running example of Algorithm 4, with parameters $r = 8$ and $\rho = 8$.

[1] F. N. Afrati, P. Koutris, D. Suciu, and J. D. Ullman. Parallel skyline queries. *Theory Comput. Syst.*, 57(4):1008–1037, 2015.

[2] W. Balke, U. Güntzer, and J. X. Zheng. Efficient distributed skylining for web information systems. In *EDBT*, pages 256–273, 2004.

[3] Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004.

[4] S. Börzsönyi, D. Kossmann, and K. Stocker. The skyline operator. In *ICDE*, pages 421–430, 2001.

[5] M. Braverman, F. Ellen, R. Oshman, T. Pitassi, and V. Vaikuntanathan. A tight bound for set disjointness in the message-passing model. In *FOCS*, pages 668–677, 2013.

[6] J. Chomicki, P. Ciaccia, and N. Meneghetti. Skyline queries, front and back. *SIGMOD Record*, 42(3):6–18, 2013.

[7] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. *J. Comput. Syst. Sci.*, 66(4):614–656, 2003.

[8] K. Hose and A. Vlachou. A survey of skyline processing in highly distributed environments. *VLDB J.*, 21(3):359–384, 2012.

[9] H. Köhler, J. Yang, and X. Zhou. Efficient parallel skyline processing using hyperplane projections. In *SIGMOD*, pages 85–96, 2011.

[10] H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, 1975.

[11] M. Lichman. UCI machine learning repository, 2013.

[12] E. Lo, K. Y. Yip, K. Lin, and D. W. Cheung. Progressive skylining over web-accessible databases. *Data Knowl. Eng.*, 57(2):122–147, 2006.

[13] K. Mullesgaard, J. L. Pederseny, H. Lu, and Y. Zhou. Efficient skyline computation in mapreduce. In *EDBT*, pages 37–48, 2014.

[14] D. Pertesis and C. Doulkeridis. Efficient skyline query processing in spatialhadoop. *Inf. Syst.*, 54:325–335, 2015.

[15] J. B. Rocha-Junior, A. Vlachou, C. Doulkeridis, and K. Nørvåg. Agids: A grid-based strategy for distributed skyline query processing. In *Data Management in Grid and Peer-to-Peer Systems, Second International Conference*, pages 12–23, 2009.

[16] G. Trimponias, I. Bartolini, D. Papadias, and Y. Yang. Skyline processing on distributed vertical decompositions. *IEEE Trans. Knowl. Data Eng.*, 25(4):850–862, 2013.

[17] A. Vlachou, C. Doulkeridis, and Y. Kotidis. Angle-based space partitioning for efficient parallel skyline computation. In *SIGMOD*, pages 227–238, 2008.

[18] S. Wang, B. C. Ooi, A. K. H. Tung, and L. Xu. Efficient skyline query processing on peer-to-peer networks. In *ICDE*, pages 1126–1135, 2007.

[19] P. Wu, C. Zhang, Y. Feng, B. Y. Zhao, D. Agrawal, and A. El Abbadi. Parallelizing skyline queries for scalable distribution. In *EDBT*, pages 112–130, 2006.

[20] L. Zhu, Y. Tao, and S. Zhou. Distributed skyline retrieval with low bandwidth consumption. *IEEE Trans. Knowl. Data Eng.*, 21(3):384–400, 2009.

APPENDIX

A. A RUNNING EXAMPLE FOR THE INTERACTIVE PRUNING ALGORITHM

Figure 16 shows a running example of Algorithm 4 with parameters $r = 8$ and $\rho = 8$. In the first round, Alice sends G_{x1} and Bob sends G_{y1} to the coordinator. In the second round, the coordinator sends the IDs of points in G_{x1} and G_{y1} to Bob and Alice respectively, and then Bob and Alice send the y-coordinates and x-coordinates of points in G_{x1} and G_{y1} respectively to the coordinator. The coordinator recovers the points in G_{x1} and G_{y1} and adds them to R (the set of candidate skyline points), and updates $l_x = 8$, $l_y = 7$, and $f_x = f_y = 2$. In the third and fourth rounds, the coordinator chooses to recover points in G_{y2} because $l_y - f_y < l_x - f_x$, and then updates $l_x = 7$. In the fifth and sixth rounds, the coordinator chooses to recovers points in G_{y3} because $l_y - f_y < l_x - f_x$, and then updates $l_x = 4$. In the last two rounds, the coordinator asks Alice for the information of the points in the remaining groups G_{x2} and G_{x3} (i.e., the groups that have *not* been recovered or pruned) since $l_x - f_x < l_y - f_y$.